



Project no.: PIRSES-GA-2011-295261
Project full title: Mobility between Europe and Argentina applying Logics to Systems
Project Acronym: MEALS
Deliverable no.: 3.2 / 3
Title of Deliverable: Generalized differential privacy: regions of priors that admit robust optimal mechanisms

Contractual Date of Delivery to the CEC:	1-Apr-2015
Actual Date of Delivery to the CEC:	15-Mar-2015
Organisation name of lead contractor for this deliverable:	INR
Author(s):	Ehab ElSalamouny Konstantinos Chatzikokolakis Catuscia Palamidessi
Participants(s):	INR
Work package contributing to the deliverable:	WP3
Nature:	R
Dissemination Level:	Public
Total number of pages:	29
Start date of project:	1 Oct. 2011 Duration: 48 month

Abstract:

Differential privacy is a notion of privacy that was initially designed for statistical databases, and has been recently extended to a more general class of domains. Both differential privacy and its generalized version can be achieved by adding random noise to the reported data. Thus, privacy is obtained at the cost of reducing the data's accuracy, and therefore their *utility*.

In this paper we consider the problem of identifying *optimal* mechanisms for generalized differential privacy, i.e. mechanisms that maximize the utility for a given level of privacy. The utility usually depends on a prior distribution of the data, and naturally it would be desirable to design mechanisms that are *universally optimal*, i.e., optimal for all priors. However it is already known that such mechanisms do not exist in general. We then characterize maximal *classes of priors* for which a mechanism which is optimal for all the priors of the class *does exist*. We show that such classes can be defined as convex polytopes in the priors space.

As an application, we consider the problem of privacy that arises when using, for instance, location-based services, and we show how to define mechanisms that maximize the quality of service while preserving the desired level of geo-indistinguishability.

Contents

1	Prologue	3
2	Introduction	3
3	Preliminaries	6
3.1	Generalized privacy	6
3.2	Differential privacy	8
3.3	Geo-indistinguishability	9
3.4	Utility model	9
3.5	Min-mutual information	10
4	Regular priors	11
5	Upper bounds for utility and min-mutual information	13
5.1	Utility	14
5.2	Min-mutual information	17
5.3	Quantifying the leakage about the database	18
6	Tight-constraints mechanisms	19
6.1	Properties	20
6.2	Tight-constraints mechanism for symmetric metric spaces	23
7	Case-studies	24
7.1	Differential privacy: sum and 2-count queries	24
7.2	Geo-indistinguishability	25
8	Conclusion and future work	26
	Bibliography	27
	MEALS Partner Abbreviations	29

1 Prologue

Privacy is an instance of the general problem of information protection, which constitutes one of the main topics of the research of our team Com ete. The history of our interest for this topic has an important milestone in the visit of Prakash to Com ete in 2006, in the context of our  quipe associ ee Printemps. We had been working for a while on a probabilistic approach to anonymity, and when Prakash arrived, he suggested to consider an information-theoretic approach instead. This was the beginning of a very fruitful collaboration between Prakash and our team, and two of the papers that originated from this collaboration became the backbone of the PhD thesis of Konstantinos Chatzikokolakis. Furthermore, the collaboration with Prakash influences, still today, our research on information protection, in the sense that our research is characterized by the paradigmatic view of a system as a noisy channel – the central concept of information theory. The present paper, which explores the properties of the channel matrix in the context of differential privacy, is a tribute to the fundamental role that Prakash has had in Com ete’s scientific life and evolution.

2 Introduction

It is often the case that a privacy threat arises not because of direct access to sensitive data by unauthorized agents, but rather because of the information they can infer from correlated public data. This phenomenon, known as *information leakage*, is quite general and it has been studied in several different domains, including programming languages, anonymity protocols, and statistical databases (see, for instance, [1, 2, 3]). Naturally, the settings and the approaches vary from domain to domain, but the principles are the same.

In the case of statistical databases, the public information is typically defined by the kind of queries we are allowed to ask, and the concerns for privacy focus on the consequences that the participation in the databases may have for the confidential data of a *single individual*. Differential privacy [4, 5] was designed to control these consequences. Since it has been recognized that the deterministic methods offer little resistance to composition attacks (i.e. to the combination of information inferred from different databases, see for instance [6, 7]), differential privacy targets probabilistic mechanisms, i.e. mechanisms that answer the query in a probabilistic fashion. Typically, they generate the output by adding random noise to the true answer, according to some probabilistic distribution. The aim of differential privacy is to guarantee that the participation of a single individual in the database will not affect too much the probability of each reported answer. More precisely, (the log of) the ratio between the likelihoods of obtaining a certain answer, from any two *adjacent* databases (i.e., differing only for the presence of an individual), must not exceed a given parameter ϵ . The rationale of this notion comes from the fact that it is equivalent to the property that the reported answer does not change significantly the probabilistic knowledge of the individual data. Differential privacy has become very popular thanks to the fact that it is easy to implement: it is sufficient to add Laplacian noise to the true answer. Furthermore, the notion and the implementation are independent from the side knowledge of the adversary about the underlying database (represented as a prior probability distribution over possible databases).

Finally, it is compositional, in the sense that the privacy loss caused by the combination of attacks is the sum of the single privacy losses.

There have been several studies aimed at applying differential privacy to other areas. In this work, we focus on the approach proposed in [8], which introduced the concept of d_X -privacy, suitable for any domain \mathcal{X} equipped with a notion of distance d_X . Given a mechanism K from the set of secrets \mathcal{X} to distribution over some set of outputs \mathcal{Z} , we say that K satisfies d_X -privacy if for any two secrets x_1 and x_2 , and any output z , the log of the ratio between $K(x_1)$ and $K(x_2)$ does not exceed $d_X(x_1, x_2)$. Note that d_X -privacy is an extension of differential privacy: the latter can be obtained by setting \mathcal{X} to be the set of databases (seen as tuples of individual records) and d_X to be the Hamming distance between these tuples, scaled by ϵ . Furthermore, it is a *conservative extension*, in the sense that it preserves the implementability by means of Laplacian noise, the independence from the prior probability, the interpretation in terms of probabilistic knowledge, and the compositionality properties. From the practical point of view, d_X -privacy is particularly suitable to protect the accuracy of the values, like in the case of smart-meter signatures [8] and the precise geographical position in location-based services [9]. Similar extensions of differential privacy obtained by generalizing the distance or the adjacency relation have been considered in [10, 11, 12].

Besides guaranteeing privacy, a mechanism should of course provide an answer which is “useful” enough for the service it has been designed. This second goal is measured in terms of *utility*, which represents the average gain that a rational user obtains from the reported answer. More precisely, let y be the true answer and let z be the output reported by the mechanism. On the basis of the latter, the user tries to make a guess y' (remapping) about the (hidden) true answer y . His gain $g(y, y')$ is determined by a given function g . The utility is then defined as *the expected gain under the best possible remapping*. While the gain function can take various forms, in this paper we restrict our analysis to the *binary* gain function, which evaluates to 1 when the user’s guess is the same as the query result ($y = y'$) and evaluates to 0 otherwise.

Obviously, there is a trade-off between privacy and utility, and we are interested in mechanisms that offer maximal utility for the desired level of d_X -privacy. Such mechanisms are called *optimal*. Naturally, we are also interested in mechanisms that are *universally* optimal, i.e., optimal under any prior¹, as we don’t want to design a different mechanism for each user². A famous result by Gosh et al. [13] states that this is possible for the *counting queries*, namely the queries of the form “how many records in the database have the property p ”, for some p . Unfortunately Brenner and Nissim showed that in differential privacy universally optimal mechanisms do not exist for any other kind of query [14]. However, one can still hope that it is possible to design mechanisms that are optimal for a significant class of users. These are exactly the main objectives of this paper: identify regions of priors which admit a *robust* optimal mechanism, i.e. a mechanism whose optimality is not affected by changes in the prior (within the region), and provide a method to construct such mechanism.

A related issue that we consider in this paper is the amount of information leaked by a mechanism, a central concept in the area of *quantitative information flow*. There have been various

¹Note that, in contrast to d_X -privacy, utility *does* depend on the prior.

²We recall that the prior represents the side knowledge of the user.

proposals for quantifying the information leakage, we consider here an *information-theoretic approach* based on Rényi min-entropy [15, 16], which is suitable for one-try attacks. A main difference between the min-entropy leakage and d_x -privacy is that the former measures the *expected* risk of disclosure of sensitive information, while the latter focuses on the worst case, i.e., it considers catastrophic any such disclosure, no matter how unlikely it is.

Recently, researchers have investigated the relation between differential privacy and min-entropy leakage [17, 18, 19], and in particular it has been proved in [18] that differential privacy induces a bound on the min-entropy leakage, which is met by a certain mechanism for the uniform prior (for which min-entropy leakage is always maximum). In this paper, we extend the above result to provide a more accurate bound for any prior in the special regions described above. More precisely, we provide a bound to the leakage specific to the prior and that can be met, under a certain condition, by a suitable mechanism.

Contributions

- We identify, for an arbitrary metric space (\mathcal{Y}, d_y) , the class of the d_y -regular distributions of \mathcal{Y} . The interest of this class is that for each prior distribution in it we are able to provide a specific upper bound to the utility of any d_y -private mechanism. We characterize this class as a geometric region, and we study its properties.
- We describe a d_y -private mechanism, called “tight-constraints mechanism”, which meets the upper bound for every d_y -regular prior, and is therefore robustly optimal in that region. We provide necessary and sufficient conditions for the existence of such mechanism, and an effective method to test the conditions and to construct the mechanism.
- We consider the domain of databases (\mathcal{X}, d_x) , where d_x is the Hamming distance, and we recast the above definitions and results in terms of min-entropy leakage. We are able to improve a result from the literature which says that differential privacy induces a bound on the min-entropy leakage for the uniform prior: We provide more accurate bounds, and show that these bounds are valid for all the d_x -regular priors (not just for the uniform one). A construction similar to the one in the previous point yields the tight-constraints mechanism which reaches those upper bounds.

A preliminary version of this paper, restricted to standard differential privacy, and without proofs, appeared in POST 2013.

Plan of the paper In the next section we recall the basic definitions of generalized differential privacy, utility, and min-entropy mutual information. Section 4 introduces the notion of d_y -regular prior, investigates the properties of these priors, and gives a geometric characterization of their region. Section 5 shows that for all d_y -regular priors on the true answers (resp. databases), d_y -privacy induces an upper bound on the utility (resp. on the min-entropy leakage). Section 6 identifies a mechanism which reaches the above bounds for every d_y -regular prior, and that is therefore the universally optimal mechanism (resp. the maximally leaking mechanism) in the region. Section 7 illustrates our methodology and results using the example of the sum queries and location privacy. Section 8 concludes and proposes some directions for future research.

3 Preliminaries

In this section we recall the generalized variant of differential privacy from [8], considering an arbitrary set of secrets \mathcal{X} , equipped with a metric $d_{\mathcal{X}}$. We then discuss two instantiations of the general definition: first, *standard differential privacy* is defined on databases under the Hamming distance. Second, *geo-indistinguishability* [9], a notion of location privacy, is obtained by using geographical locations as secrets, under the Euclidean distance. Finally, we recall a standard way for measuring the utility of a mechanism, and the notion of min-mutual information.

3.1 Generalized privacy

As discussed in the introduction, a generalized variant of differential privacy can be defined on an arbitrary set of secrets \mathcal{X} , equipped with a metric $d_{\mathcal{X}}$. Intuitively, $d_{\mathcal{X}}(x, x')$ gives the “distinguishability level” between secrets x, x' , based on the privacy semantics that we wish to obtain. The smaller the distinguishability level is, the harder it should be for the adversary to distinguish the two secrets, hence offering privacy, while secrets at great distance are allowed to be distinguished, giving the possibility to obtain some controlled knowledge about the secret.

A *mechanism* from \mathcal{X} to \mathcal{Z} is a function $K : \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Z})$, where $\mathcal{P}(\mathcal{Z})$ denotes the set of probability distributions over some set of outputs \mathcal{Z} . In this paper we consider \mathcal{X}, \mathcal{Z} to be finite, hence the involved distributions to be discrete. The mechanism’s outcome $K(x)$ is then a probability distribution, and $K(x)(z)$ is the probability of an output $z \in \mathcal{Z}$ when running the mechanism on $x \in \mathcal{X}$. For simplicity we write $K : \mathcal{X} \rightarrow \mathcal{Z}$ to denote a mechanism from \mathcal{X} to \mathcal{Z} (omitting \mathcal{P}).

The multiplicative distance $d_{\mathcal{P}}$ between probability distributions $\mu_1, \mu_2 \in \mathcal{P}(\mathcal{Z})$ is defined as $d_{\mathcal{P}}(\mu_1, \mu_2) = \sup_{z \in \mathcal{Z}} |\ln \frac{\mu_1(z)}{\mu_2(z)}|$ with the convention that $|\ln \frac{\mu_1(z)}{\mu_2(z)}| = 0$ if both $\mu_1(z), \mu_2(z)$ are zero and ∞ if only one of them is zero.

We are now ready to give the definition of $d_{\mathcal{X}}$ -privacy:

Definition 1. A mechanism $K : \mathcal{X} \rightarrow \mathcal{Z}$ satisfies $d_{\mathcal{X}}$ -privacy, iff $\forall x, x' \in \mathcal{X}$:

$$d_{\mathcal{P}}(K(x), K(x')) \leq d_{\mathcal{X}}(x, x')$$

or equivalently:

$$K(x)(z) \leq e^{d_{\mathcal{X}}(x, x')} K(x')(z) \quad \forall z \in \mathcal{Z}$$

The intuition behind this definition is that the attacker’s ability to distinguish two secrets should depend on their distinguishability level $d_{\mathcal{X}}(x, x')$. The closer two secrets are, the more similar the mechanism’s output on those secrets should be, making it harder for the adversary to distinguish them. Depending on the choice of $d_{\mathcal{X}}$, the definition can be adapted to the application at hand, giving rise to different notions of privacy.

In [8], two alternative characterizations of $d_{\mathcal{X}}$ -privacy are also given, in which the attacker’s knowledge is explicitly quantified, which makes it easier to understand the privacy guarantees obtained by a particular choice of $d_{\mathcal{X}}$.

Answering queries. In practice, we often want to learn some information about our secret, that is we want to obtain the answer to a query $f : \mathcal{X} \rightarrow \mathcal{Y}$. To do so privately, we can compose f with a “noise” mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$, thus obtaining an “oblivious” mechanism $H \circ f : \mathcal{X} \rightarrow \mathcal{Z}$, called oblivious since the answer depends only on $f(x)$ and not on x itself. The role of H is to add random noise to the true query result $f(x)$ and produce a “noisy” reported output $z \in \mathcal{Z}$.

Since we assume all sets to be finite, the mechanism H can be described by a stochastic matrix $H = (h_{yz})$, called the *noise matrix*, whose rows are indexed by the elements of \mathcal{Y} and whose columns are indexed by the elements of \mathcal{Z} . Hence, h_{yz} is the probability of reporting z when the true query result is y .

Given a metric d_y on \mathcal{Y} , the generalized definition of privacy allows us to directly talk about the privacy of H , without involving f at all. Using matrix notation, d_y -privacy for H (Definition 1) can be written as

$$h_{yz} \leq e^{d_y(y,y')} h_{y'z} \quad \forall y, y' \in \mathcal{Y}, z \in \mathcal{Z} \quad (1)$$

A natural question, then, is how d_x -privacy of the composed mechanism $H \circ f$ relates to d_y -privacy of H . The connection between the two comes from the concept of *uniform Δ -sensitivity*.

Definition 2. A sequence y_1, \dots, y_n is called a *chain* from y_1 to y_n . We say that such chain is *tight* if $d_y(y_1, y_n) = \sum_i d_y(y_i, y_{i+1})$. Two elements $y, y' \in \mathcal{Y}$ are called Δ -*expansive* iff $d_y(y, y') = \Delta d_x(x, x')$ for some $x \in f^{-1}(y), x' \in f^{-1}(y')$. A chain is Δ -*expansive* iff all steps y_i, y_{i+1} are Δ -*expansive*.

Finally, f is uniformly Δ -sensitive wrt d_x, d_y iff:

- for all $x, x' \in \mathcal{X}$: $d_y(f(x), f(x')) \leq \Delta d_x(x, x')$, and
- for all $y, y' \in \mathcal{Y}$: there exists a tight and Δ -expansive chain from y to y' .

The intuition behind this definition is that f expands distances by at most Δ , and there are no answers that are always the results of a smaller expansion: all $y, y' \in \mathcal{Y}$ can be linked by a chain in which the expansion is exactly Δ . Under this condition, it has been shown in [8] that the privacy of H characterizes that of $H \circ f$.

Theorem 1 ([8]). *Assume that f is uniformly Δ -sensitive wrt d_x, d_y . Then H satisfies d_y -privacy if and only if $H \circ f$ satisfies Δd_x -privacy.*

In the remaining of the paper, we give results about d_y -privacy for H , for an arbitrary metric d_y , independently from any function f . The results can be used either to talk about the privacy of H itself, or – given the above theorem – about the privacy of oblivious mechanisms of the form $H \circ f$, for some function f for which uniform sensitivity can be established. A typical case of uniform sensitivity arises in standard differential privacy when d_y is the metric obtained from the *induced graph* of f , as discussed in the next section. But uniform sensitivity can be established for other types of metrics; some examples are given in [8].

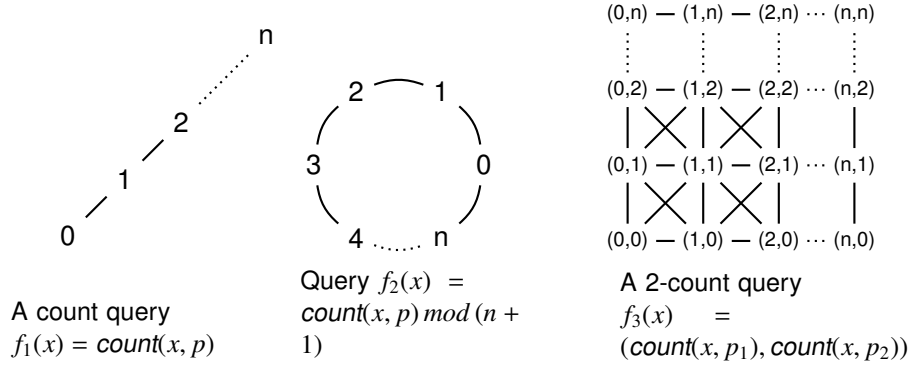


Figure 1: The induced graph of different queries

3.2 Differential privacy

The notion of differential privacy, introduced by Dwork in [4], imposes constraints on data reporting mechanisms so that the outputs produced by two databases differing only for one record are almost indistinguishable. Let V be a universe of values and u the number of individuals. The set of all possible databases (u -tuples of values from V) is $\mathcal{V} = V^u$. Two databases $x, x' \in \mathcal{V}$ are called *adjacent*, written $x \sim x'$, iff they differ in the value of exactly one individual. The adjacency relation \sim defines a graph, and the length of the shortest path between two databases x, x' in the graph, written $d_h(x, x')$, defines a metric called the Hamming distance. In other words, $d_h(x, x')$ is the number of individuals in which x and x' differ.

The property of ϵ -differential privacy requires that, for any two adjacent databases, the ratio of the probabilities of producing a certain output is bound by e^ϵ . It is easy to see that this property is equivalent to ϵd_h -privacy, under the Hamming distance d_h .

Given a query $f : \mathcal{V} \rightarrow \mathcal{Y}$, the adjacency relation \sim can be extended to \mathcal{Y} , giving rise to the *induced graph* \sim_f of f [14, 19], defined as:

$$y \sim_f y' \quad \text{iff} \quad x \sim x' \text{ for some } x \in f^{-1}(y), x' \in f^{-1}(y')$$

Figure 1 shows the induced graph of three different queries. In these examples $\text{count}(x, p)$ refers to a counting query which returns the number of records in the database x which satisfy a certain property p . Other queries in the figure are expressed using the *count* function.

Furthermore, let $d_{\sim_f}(y, y')$ be the metric on \mathcal{Y} defined as the shortest \sim_f -path from y to y' . It has then been shown in [8] that any function f is uniformly 1-sensitive wrt d_h, d_{\sim_f} . As a consequence of this, and of Theorem 1, ϵ -differential privacy of an oblivious mechanism $H \circ f$ can be characterized by the ϵd_{\sim_f} -privacy of H .

Corollary 1. *For any query $f : \mathcal{V} \rightarrow \mathcal{Y}$, H satisfies ϵd_{\sim_f} -privacy if and only if $H \circ f$ satisfies ϵd_h -privacy.*

3.3 Geo-indistinguishability

An advantage of the generalized definition of privacy is that it can be applied in cases when there is a single individual involved – hence the notion of adjacency is inadequate – by using a metric that gives a meaningful notion of privacy for the application at hand. An example of such a notion is *geo-indistinguishability* [9], proposed as a formal notion of location privacy in the context of Location Based Services (LBSs).

Consider a mobile user, typically using a GPS-enabled hand-held device, who wishes to obtain information related to his current location, for instance restaurants close to him. To do so, he can query an LBS provider, providing his actual location x as part of the query. However, location information is not only inherently sensitive itself, but also correlated to a variety of other sensitive information, such as political and religious beliefs, medical information, etc. Hence, the user would like to perform the LBS query privately, that is without disclosing his exact location to the provider. Note that protecting the user’s *identity* is not the goal here; in fact, the user might wish to be authenticated to the service provider in order to obtain personalized recommendations. What he is interested in, instead, is hiding his *location*.

A possible solution is to use a *location obfuscation* mechanism [20], producing a noisy location z which is reported to the service provider. A natural goal then is to formalize the privacy guarantees provided by such a mechanism, for which various approaches have been proposed in the literature [21].

Geo-indistinguishability provides such a formal definition of location privacy, and can be expressed as an instance of d_x -privacy. Secrets \mathcal{X} are now locations (a subset of \mathbb{R}^2), and ϵ -geo-indistinguishability is ϵd_2 -privacy, where d_2 is the Euclidean distance between locations.³ Intuitively, $d_{\mathcal{P}}(K(x), K(x')) \leq \epsilon d_2(x, x')$ requires that the closer (geographically) two locations x, x' are, the more likely to produce the same reported location z they should be. This allows the provider to get some approximate information necessary to provide the service (e.g. distinguish locations in Paris from those in London), but prevents him from learning x with high accuracy (since locations x' close to x produce the same z with similar probabilities).

The results of this paper refer to an arbitrary metric between secrets, hence they are directly applicable to geo-indistinguishability. A case-study in the context of location privacy is given in Section 7.2.

3.4 Utility model

The main role of a noise mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ is to guarantee d_y -privacy while providing useful information about the true query result, i.e. to satisfy a trade-off between the privacy and utility. For quantifying the utility of H we follow a standard model from [13]. Let $y \in \mathcal{Y}$ be the result of executing a query f . The mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ processes y and produces an output z in some domain \mathcal{Z} to the user. Based on the reported output z and prior knowledge about the likely results of f , she applies a remapping function $R : \mathcal{Z} \rightarrow \mathcal{Y}$ to z to produce a guess $y' \in \mathcal{Y}$ for the real query result. Note that the composite mechanism $R \circ H : \mathcal{Y} \rightarrow \mathcal{Y}$ is a mechanism

³Note that any other meaningful geographical distance could also be used, such as the Manhattan or a map-based distance.

whose output domain is the query results domain \mathcal{Y} . We say that H is remapped to $R \circ H$ by the remap R . Now, with the user's guessed value y' , a real-valued *gain function* $g : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ quantifies how informative y' is compared to the real query result y . In this paper we restrict our analysis to the binary gain function g_b which is defined as $g_b(y, y') = 1$ iff $y' = y$ and 0 otherwise. The choice of this gain corresponds to the preference of a user to guess the true query result.

In practice, the user usually bases her guess y' about the real query result on prior knowledge about the underlying secret and the underlying query. This knowledge is modeled by a probability distribution π (called *prior*) over the domain \mathcal{Y} of query results. Now the utility of a mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ with respect to a prior π and a remap $R : \mathcal{Z} \rightarrow \mathcal{Y}$ is the expected value of the underlying gain function g_b , and is therefore expressed as

$$\mathcal{U}(H, \pi, R) = \sum_{y, y'} \pi_y (HR)_{yy'} g_b(y, y'). \quad (2)$$

Using the definition of g_b , the above expression reduces to a convex combination of the diagonal elements of HR as follows.

$$\mathcal{U}(H, \pi, R) = \sum_y \pi_y (HR)_{yy}. \quad (3)$$

Accordingly, we say that a d_y -private mechanism H is d_y -optimal for a prior π if there is a remap R such that $\mathcal{U}(H, \pi, R)$ is maximal for all d_y -private mechanisms and all remaps.⁴ In general the optimality of a mechanism depends on the prior (related to the user). That is a mechanism that is optimal for a prior may not be optimal for another one. In the setting of differential privacy, it has been proven [14] that for any query, other than a single counting one, there is no mechanism that is optimal for all priors simultaneously. Nevertheless, we identify in Section 4 a region of priors, where it is possible to find a single mechanism which is optimal to all of them.

3.5 Min-mutual information

In this section we recall the use of an information-theoretic notion, namely mutual information, to quantify the amount of information conveyed by a mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ as an information theoretic channel.

Following recent works in the area of quantitative information flow ([15, 16, 17]), we adopt Rényi's *min-entropy* ([22]) as our measure of uncertainty. The min-entropy $\mathcal{H}_\infty(\pi)$ of a prior π , defined as $\mathcal{H}_\infty(\pi) = -\log_2 \max_i \pi_i$, measures the user's uncertainty about the query result. Then, the corresponding notion of *conditional min-entropy*, defined as $\mathcal{H}_\infty(H, \pi) = -\log_2 \sum_{z \in \mathcal{Z}} \max_y \pi_y h_{yz}$, measures the uncertainty about the query result after observing an output $z \in \mathcal{Z}$. Finally, subtracting the latter from the former brings us to the notion of min-mutual information:

$$\mathcal{L}(H, \pi) = \mathcal{H}_\infty(\pi) - \mathcal{H}_\infty(H, \pi)$$

which measures the amount of information about the query result conveyed by the mechanism H . In the area of quantitative information flow this quantity is known as *min-entropy leakage*; the reader is referred to [15] for more details about this notion.

⁴Note that there may exist many optimal mechanisms for a given prior.

4 Regular priors

In this section we describe a region of priors, called ‘ d_y -regular’. These priors are determined by the metric d_y on the domain \mathcal{Y} . Recall that the d_y -privacy constraints for H can be written as $h_{yz}/h_{y'z} \geq e^{-d_y(y,y')}$ for all $y, y' \in \mathcal{Y}$. Since every lower bound $e^{-d_y(y,y')}$ depends only on y, y' , the constraints can be described altogether by a square matrix Φ formed by such lower bounds. We refer to this matrix as the *privacy-constraints* matrix.

Definition 3 (privacy-constraints matrix). The *privacy-constraints matrix* Φ of a metric d_y is a square matrix, indexed by $\mathcal{Y} \times \mathcal{Y}$, where $\phi_{yy'} = e^{-d_y(y,y')}$ for all $y, y' \in \mathcal{Y}$.

Note that Φ is symmetric ($\phi_{yy'} = \phi_{y'y}$) due to the symmetry of d_y . Recall that d_y describes the privacy restrictions imposed on the domain \mathcal{Y} . In particular these restrictions become vacuous if $d_y(y, y') \rightarrow \infty$ for all $y, y' : y \neq y'$. In this extreme case the privacy-constraints matrix Φ converges to the identity matrix where each diagonal entry is 1 and all other entries are 0. We now define the d_y -regular priors, in terms of the privacy-constraints matrix of d_y . For a vector μ having cardinality $|\mathcal{Y}|$, we use $\mu \geq \mathbf{0}$ to denote $\forall y : \mu_y \geq 0$.

Definition 4 (d_y -regular prior). A prior π is called *d_y -regular* iff there exists a row vector $\mu \geq \mathbf{0}$ such that $\pi = \mu \Phi$.

In the following we describe the common properties of these priors and also give a geometric characterization for their region comparing it to the whole prior space. As a first observation, this region converges to the entire prior space when the privacy constraints on \mathcal{Y} become vacuous. This is because, as described above, Φ approaches the identity matrix where the vector μ exists for each prior π (just define $\mu = \pi$).

An important property of any d_y -regular prior is that the ratio between any two of its entries $\pi_y, \pi_{y'}$ is always bound by $e^{d_y(y,y')}$. Because of this property, such a prior is called *d_y -regular*.

Proposition 1. For every d_y -regular prior π and for all $y, y' \in \mathcal{Y}$ we have that $\pi_y/\pi_{y'} \leq e^{d_y(y,y')}$.

Proof. By Definition 4, the ratio $\pi_y/\pi_{y'}$ is given by

$$\pi_y/\pi_{y'} = \frac{\sum_{y''} \mu_{y''} \phi_{y''y}}{\sum_{y''} \mu_{y''} \phi_{y''y'}}. \quad (4)$$

By the definitions of $\phi_{y''y}, \phi_{y''y'}$ we also have that

$$\phi_{y''y'} = e^{-d_y(y'',y')} \geq e^{-(d_y(y'',y) + d_y(y,y'))} = e^{-d_y(y,y')} \phi_{y''y}.$$

The above inequality is implied by the triangle inequality, $d_y(y'', y') \leq d_y(y'', y) + d_y(y, y')$ and the fact that $e^{-1} < 1$. Since $\mu_{y''} \geq 0$ for all y'' , we have

$$\sum_{y''} \mu_{y''} \phi_{y''y'} \geq e^{-d_y(y,y')} \sum_{y''} \mu_{y''} \phi_{y''y}$$

Substituting the above inequality in Eq. (4) completes the proof. □ □

The above property restricts nearby elements of \mathcal{Y} (with respect to the metric d_y) to have ‘similar’ probabilities. In practice, this property holds for a large class of users who have no sharp information that discriminates between nearby elements of \mathcal{Y} . Note that the above property is not equivalent to Definition 4. Namely, it is not true that all priors having such a property are d_y -regular.

A consequence of the above proposition is that for any d_y -regular prior π , the probability π_y associated with $y \in \mathcal{Y}$ is restricted by upper and lower bounds as follows.

Proposition 2. *For every d_y -regular prior π and for every $y \in \mathcal{Y}$ we have that*

$$1 / \sum_{y' \in \mathcal{Y}} e^{d_y(y, y')} \leq \pi_y \leq 1 / \sum_{y' \in \mathcal{Y}} e^{-d_y(y, y')}.$$

Proof. By Proposition 1, it holds for every pair of entries $\pi_y, \pi_{y'}$ that

$$\pi_{y'} \leq e^{d_y(y, y')} \pi_y \quad \text{and} \quad e^{-d_y(y, y')} \pi_y \leq \pi_{y'}.$$

Summing the above inequalities over y' , we get

$$\sum_{y' \in \mathcal{Y}} \pi_{y'} \leq \pi_y \sum_{y' \in \mathcal{Y}} e^{d_y(y, y')} \quad \text{and} \quad \pi_y \sum_{y' \in \mathcal{Y}} e^{-d_y(y, y')} \leq \sum_{y' \in \mathcal{Y}} \pi_{y'}.$$

Since $\sum_{y' \in \mathcal{Y}} \pi_{y'} = 1$, the above inequalities imply the upper and lower bounds for π_y . \square \square

One obvious implication is that any d_y -regular prior must have full support, that is $\pi_y > 0$ for all $y \in \mathcal{Y}$. In the following we describe the set of d_y -regular priors as a region in the prior space. For doing so, we first define in the following set of priors which we refer to as the corner priors.

Definition 5 (corner priors). For every $y \in \mathcal{Y}$, a corresponding corner prior, denoted by \mathbf{c}^y , is defined as

$$c_{y'}^y = \frac{\phi_{yy'}}{\sum_{y'' \in \mathcal{Y}} \phi_{yy''}} \quad \forall y' \in \mathcal{Y}.$$

Note that the above definition is sound, i.e. \mathbf{c}^y is a probability distribution for all $y \in \mathcal{Y}$. Note also that there are $|\mathcal{Y}|$ corner priors; each one corresponds to an element $y \in \mathcal{Y}$. By inspecting the entries of \mathbf{c}^y , observe that c_y^y has the maximum value compared to other entries, and moreover this value is exactly the upper bound specified by Proposition 2. We can therefore interpret this observation informally as \mathbf{c}^y is ‘maximally biased’ to y . It can be also seen that each corner prior is d_y -regular. In fact for any corner \mathbf{c}^y , there is a row vector $\boldsymbol{\mu}$ that satisfies the condition in Def. 4; this vector is obtained by setting $\mu_y = 1 / \sum_{y' \in \mathcal{Y}} \phi_{yy'}$ and $\mu_{y'} = 0$ for all $y' \neq y$. Here it is easy to verify that $\mathbf{c}^y = \boldsymbol{\mu} \Phi$.

Now we can describe the region of the d_y -regular priors using the corner priors. Precisely, this region consists of all convex combinations of the corner priors.

Proposition 3 (convexity). *A prior π is d_y -regular iff it is a convex combination of the corner priors, i.e. there exist real numbers $\gamma_y \geq 0$, $y \in \mathcal{Y}$ such that*

$$\boldsymbol{\pi} = \sum_{y \in \mathcal{Y}} \gamma_y \mathbf{c}^y \quad \text{and} \quad \sum_{y \in \mathcal{Y}} \gamma_y = 1.$$

Proof. By Definition 4, a prior π is d_y -regular iff there exists vector $\mu \geq \mathbf{0}$ such that $\pi = \mu \Phi$; that is iff there are reals $\mu_y \geq 0$ for all $y \in \mathcal{Y}$, such that π can be written as a linear combination of Φ 's rows as follows.

$$\pi = \sum_{y \in \mathcal{Y}} \mu_y \Phi_y,$$

where Φ_y is the row of Φ corresponding to the element $y \in \mathcal{Y}$. From Def. 5, observe that each row Φ_y is equal to $(\sum_{y' \in \mathcal{Y}} \phi_{yy'}) c^y$. By substitution in the above equation for π , we get that π is d_y -regular iff $\pi = \gamma_y \Phi_y$ where $\gamma_y = \mu_y (\sum_{y' \in \mathcal{Y}} \phi_{yy'})$. Note that the existence of the vector $\mu \geq 0$ is equivalent to the existence of the coefficients $\gamma_y \geq 0$. Observe also that $\sum_{y \in \mathcal{Y}} \gamma_y = \mu \Phi = 1$. These observations complete the proof. \square \square

From Proposition 3 the region of d_y -regular priors is a convex set, where each point (prior) in this region is a convex combination of the corner priors. This region is therefore geometrically regarded as a convex polytope in the prior space. Since the corner points always exists, this region is never empty. For a prior π in this region, the coefficients γ_y model the ‘proximity’ of π to each corner prior c^y . In particular, note that $0 \leq \gamma_y \leq 1$, and $\gamma_y = 1$ iff $\pi = c^y$. We demonstrate this geometric interpretation using the following examples.

Example 1. Consider a simple domain \mathcal{Y} consisting of 3 elements organized in a graph structure where $d_g(y, y')$ is the graph distance between y, y' . Now for an arbitrary scaling number $\epsilon > 0$, we can define the metric d_y as $d_y(y, y') = \epsilon d_g(y, y')$. Since every prior on \mathcal{Y} has 3 entries (specifying the probability of every element $y \in \mathcal{Y}$), the prior space for \mathcal{Y} can be represented by the 3-dimensional Euclidean space. Figure 2 visualizes the region of d_y -regular priors in two cases: when the graph structure of \mathcal{Y} is a line, and when it is a circle. Note that in both cases, we have 3 corner priors c^1, c^2, c^3 . In each case, the region is depicted for $\epsilon = 0.7$ and $\epsilon = 1.6$. Note in this example that ϵ controls the privacy constraints imposed by d_y -privacy, which in turn determine the size of the region of d_y -regular priors. In particular with $\epsilon = 1.6$ (less privacy), the region is larger than the one with $\epsilon = 0.7$. In general the region expands as ϵ increases and converges to the entire region of priors defined by the corner points $\{(0, 0, 1), (0, 1, 0), (1, 0, 0)\}$ when $\epsilon \rightarrow \infty$.

Example 2. Suppose that \mathcal{Y} contains 4 elements, and d_y is defined as $d_y(y, y') = D$ for all $y, y' : y \neq y'$. In this case every prior contains 4 entries and therefore is not possible to be plotted in the 3-dimensional space. However, using the fact that the fourth component is redundant ($\sum_i \pi_i = 1$), every prior is fully described by its ‘projection’ onto the 3-dimensional subspace. Figure 3 shows the projection of the d_y -regular prior region for different values of D . Again the privacy constraints enforced by d_y -privacy are determined by D . The less restricted is D (i.e. having a higher value), the bigger the region is; and eventually coincides with the entire space when $D \rightarrow \infty$.

5 Upper bounds for utility and min-mutual information

In this section, we further describe the d_y -regular priors on the domain \mathcal{Y} in terms of the utility that can be achieved for these priors by a mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ satisfying d_y -privacy. We also




Figure 2: Regions of d_y -regular priors for Example 1




Figure 3: Regions of d_y -regular priors for Example 2

describe the amount of information that can be conveyed by H to users with such priors. More precisely, we identify for any d_y -regular prior π upper bounds for the utility and min-mutual information, considering all d_y -private mechanisms and all possible remaps. These bounds are indeed induced by the privacy constraints defined by the metric d_y .

5.1 Utility

For a given domain \mathcal{Y} equipped with the metric d_y , consider a d_y -private mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ producing observables in some domain \mathcal{Z} . In the following analysis we derive a linear algebraic expression for $\mathcal{U}(H, \pi, R)$, the utility of H for a prior π using the remap $R : \mathcal{Z} \rightarrow \mathcal{Y}$. Such an expression will play the main role in the subsequent results. We start by observing that the matrix product of H and the remap R describes an d_y -private mechanism $HR : \mathcal{Y} \rightarrow \mathcal{Y}$. Therefore the entries of HR satisfy the following subset of constraints.

$$e^{-d_y(y,y')} (HR)_{y'y'} \leq (HR)_{yy'}$$

for all $y, y' \in \mathcal{Y}$. Using Definition 3 of the privacy-constraints matrix Φ , and taking into account that $\sum_{y' \in \mathcal{Y}} (HR)_{yy'} = 1$ for all y (as both H and R are stochastic), we get the following inequalities.

$$\sum_{y' \in \mathcal{Y}} \phi_{yy'} (HR)_{y'y'} \leq 1, \quad \forall y \in \mathcal{Y}.$$

The inequality operators can be replaced by equalities while introducing *slack* variables $s_y : 0 \leq s_y \leq 1$ for all $y \in \mathcal{Y}$. The above inequalities can therefore be written as follows.

$$\sum_{y' \in \mathcal{Y}} \phi_{yy'} (HR)_{y'y'} + s_y = 1, \quad \forall y \in \mathcal{Y}.$$

Let the slack variables s_y form a column vector \mathbf{s} indexed by \mathcal{Y} . Let also $\mathbf{1}$ denote a column vector of the same size and having all entries equal to 1. Using these vectors and the privacy-constraints matrix Φ (for the given metric d_y), the above equations can be rewritten in the following matrix form.

$$\Phi \text{diag}(HR) + \mathbf{s} = \mathbf{1}, \quad (5)$$

where $\text{diag}(HR)$ is the column vector consisting of the diagonal entries of HR . Now, for any mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ and a remap $R : \mathcal{Z} \rightarrow \mathcal{Y}$ satisfying Eq. (5), and for a prior π , we want to refine the generic expression (3) of the utility by taking Eq. (5) into account. We start by rewriting Eq. (3) in the following matrix form.

$$\mathcal{U}(H, \pi, R) = \pi \text{diag}(HR). \quad (6)$$

Now, let μ be a row vector such that

$$\pi = \mu \Phi. \quad (7)$$

Note that, the above matrix equation is in fact a system of $|\mathcal{Y}|$ linear equations. The y th equation in this system is formed by the y th column of Φ , and the y th entry of π as follows.

$$\mu \Phi_y = \pi_y \quad \forall y \in \mathcal{Y}.$$

Solving this system of equations for the row vector μ has the following possible outcomes: If the matrix Φ is invertible, then, for any prior π , Eq. (7) has exactly one solution. If Φ is not invertible (i.e. it contains linearly dependent columns), then there are either 0 or an infinite number of solutions, depending on the prior π : If the entries of π respect the linear dependence relation then there are infinitely many solutions. Otherwise, the equations are ‘*inconsistent*’, in which case there are no solutions.

Whether Φ is invertible or not, we consider here only the priors where the matrix equation (7) has at least one solution μ . Note that, by definition, all the d_y -regular priors have this property, but there can be others for which the solution μ has some negative components. In some of the results below (in particular in Lemma 1) we consider this larger class of priors for the sake of generality.

Multiplying Equation (5) by μ yields

$$\mu \Phi \text{diag}(HR) + \mu \mathbf{s} = \mu \mathbf{1}. \quad (8)$$

Substituting Equations (7) and (6) in the above equation consecutively provides the required expression for the utility and therefore proves the following lemma.

Lemma 1. *For a metric space (\mathcal{Y}, d_y) let π be any prior on \mathcal{Y} . Then for every row vector μ satisfying $\pi = \mu \Phi$, the utility of any d_y -private mechanism H for π using a remap R is given by*

$$\mathcal{U}(H, \pi, R) = \mu \mathbf{1} - \mu s, \quad (9)$$

for a vector s satisfying $0 \leq s_y \leq 1$ for all $y \in \mathcal{Y}$.

Lemma 1 expresses the utility function for any d_y -private mechanism H , for a prior π satisfying $\pi = \mu \Phi$, and using a remap R . This utility is expressed as a function of the vector μ and the slack vector s . Although the matrix H and the remap R do not explicitly appear on the right side of Equation (9), the utility still depends on them indirectly through the vector s . Namely, according to Equation (5), the choice of H and R determines the slack vector s . The utility function depends also on the prior π , because the choice of π determines the set of vectors μ satisfying Eq. (7). Substituting any of these vectors in Eq. (9) yields the same value for $\mathcal{U}(H, \pi, R)$.

Now recall from Definition 4 that for every d_y -regular prior π there is μ satisfying $\pi = \mu \Phi$ and $\mu \geq \mathbf{0}$. This characteristic together with Lemma 1 implies an upper bound on the utility of any d_y -private mechanism H for π .

Theorem 2 (utility upper bound). *Let π be a d_y -regular prior and $H : \mathcal{Y} \rightarrow \mathcal{Z}$ be a d_y -private mechanism. Then for all row vectors $\mu \geq \mathbf{0}$ satisfying $\mu \Phi = \pi$, and any remap R , it holds that*

$$\mathcal{U}(H, \pi, R) \leq \sum_{y \in \mathcal{Y}} \mu_y. \quad (10)$$

Furthermore the mechanism H and remap R satisfy the equality in (10) for every d_y -regular prior iff $\Phi \text{diag}(HR) = \mathbf{1}$.

Proof. Since π is d_y -regular, we have $\pi = \mu \Phi$ for a vector $\mu \geq \mathbf{0}$. Applying Lemma 1 and noting that $s_y \geq 0$ for all $y \in \mathcal{Y}$, we observe that $\mu s \geq 0$ and hence the utility is upper-bounded by $\mu \mathbf{1} = \sum_{y \in \mathcal{Y}} \mu_y$.

It remains to show that this bound is attained for every d_y -regular prior if and only if $\Phi \text{diag}(HR) = \mathbf{1}$, which is equivalent (according to Eq. (5)) to $s = \mathbf{0}$: Clearly, if $s = \mathbf{0}$, then applying Lemma 1 yields the equality in (10) for every d_y -regular prior. For the ‘only if’ direction, it is sufficient to find a regular prior for which $s = \mathbf{0}$ must hold to satisfy the equality in (10). For this purpose we recall that every corner prior c^y satisfies $\mu^y \Phi = c^y$ where $\mu_y^y > 0$. Now consider the prior $\bar{\pi} = (1/|\mathcal{Y}|) \sum_{y \in \mathcal{Y}} c^y$, which is d_y -regular by Proposition 3. It is easy to see that it holds $\bar{\mu} \Phi = \bar{\pi}$ where $\bar{\mu} = (1/|\mathcal{Y}|) \sum_{y \in \mathcal{Y}} \mu^y$. Observe here that $\bar{\mu}_y > 0$ for all $y \in \mathcal{Y}$. Suppose now that the equality in (10) holds for $\bar{\mu}$. Therefore it must hold, by Lemma 1, that $\bar{\mu} s = 0$. Since $\bar{\mu}_y > 0$ for all $y \in \mathcal{Y}$, it must hold that $s = \mathbf{0}$. This completes the proof. \square \square

The above result can be also seen from the geometric perspective. As shown by Proposition 3, each member in the region of d_y -regular priors is described as a convex combination of the corner priors. That is there are coefficients $\gamma_y \geq 0$ for $y \in \mathcal{Y}$ which form this combination. It can be shown (as in the proof of Proposition 3) that $\gamma_y = \mu_y \left(\sum_{y' \in \mathcal{Y}} \phi_{yy'} \right)$. Hence, the upper bound given by Theorem 2 can be written as follows using the coefficients γ_y .

$$\mathcal{U}(H, \pi, R) \leq \sum_{y \in \mathcal{Y}} \frac{\gamma_y}{\sum_{y' \in \mathcal{Y}} \phi_{yy'}}.$$

Inspecting the above result for corner priors, recall that for a corner \mathbf{c}^y , $\gamma_{y'}$ is 1 for $y' = y$ and is 0 otherwise; thus, the utility upper bound for \mathbf{c}^y is therefore $1 / \sum_{y'} \phi_{yy'}$. Moreover, the upper bound for each d_y -regular prior $\boldsymbol{\pi}$ can be regarded (according to the above equation) as a convex combination of the upper bounds for the corner priors. That is, from the geometric perspective, the utility upper bound for $\boldsymbol{\pi}$ linearly depends on its proximity to the corner priors.

5.2 Min-mutual information

In this paper we use the information-theoretic notion of min-mutual information in two distinct ways: first, we use it to measure the information conveyed about the result of a specific query, similarly to the use of “utility” in the previous section. Mutual information and utility are indeed closely related, which allows us to transfer the bound obtained in the previous section to the information-theoretic setting.

Second, we use it to quantify the information about the secret itself, thus obtaining what is known in the area of quantitative information flow as *min-entropy leakage* [15]. The above bound can therefore be interpreted as a bound on the information leaked by any mechanism, even non-oblivious ones, independently from the actual query. For arbitrary priors, we obtain in a more natural way the bound conjectured in [17] and proven in [19]. Moreover, if we restrict to specific (d_y -regular) priors, then we are able to provide more accurate bounds.

The following result from [19] shows that min-mutual information corresponds to the notion of utility under the binary gain function and using an *optimal* remap, i.e., a remap that gives the best utility among all possible remaps, for the given prior.

Proposition 4 ([19]). *Given a mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ and a prior $\boldsymbol{\pi}$, let \hat{R} be an optimal remap for $\boldsymbol{\pi}, H$. Then, we have*

$$\mathcal{L}(H, \boldsymbol{\pi}) = \log_2 \frac{\mathcal{U}(H, \boldsymbol{\pi}, \hat{R})}{\max_y \pi_y}$$

This connection allows us to transfer the upper-bound given by Theorem 2 to min-mutual information.

Proposition 5 (min-mutual information upper bound). *Let $\boldsymbol{\pi}$ be a d_y -regular prior and $H : \mathcal{Y} \rightarrow \mathcal{Z}$ be a d_y -private mechanism. Then for all row vectors $\boldsymbol{\mu} \geq \mathbf{0}$ satisfying $\boldsymbol{\mu} \Phi = \boldsymbol{\pi}$, we have:*

$$\mathcal{L}(H, \boldsymbol{\pi}) \leq \log_2 \frac{\sum_{y \in \mathcal{Y}} \mu_y}{\max_y \pi_y}. \quad (11)$$

Furthermore, H satisfies the equality for every d_y -regular prior iff there is a remap R such that $\Phi \text{diag}(HR) = \mathbf{1}$.

Proof. By Proposition 4, the leakage $\mathcal{L}(H, \boldsymbol{\pi})$ is monotonically increasing with the utility $\mathcal{U}(H, \boldsymbol{\pi}, \hat{R})$. By Theorem 2, this utility is upper-bounded by $\sum_{y \in \mathcal{Y}} \mu_y$. Substituting this upper bound in Proposition 4 yields the inequality (11) where the equality holds iff it holds in Theorem 2 for H and an optimal remap \hat{R} . That is iff $\Phi \text{diag}(H\hat{R}) = \mathbf{1}$. This condition is equivalent to the condition of equality in Proposition 5, because if a remap R satisfies this latter condition then it must be

optimal because the utility with R (by Theorem 2) is globally maximum, that is no other remap can achieve higher utility. \square \square

The above bound holds only for d_y -regular priors. However, it is well-known ([16]) that min-mutual information is maximized by the uniform prior \mathbf{u} , i.e. $\mathcal{L}(H, \boldsymbol{\pi}) \leq \mathcal{L}(H, \mathbf{u})$ for all $H, \boldsymbol{\pi}$. Thus, in cases when \mathbf{u} is d_y -regular, we can extend the above bound to *any* prior.

Corollary 2. *Suppose that the uniform prior \mathbf{u} is d_y -regular, and let $H : \mathcal{Y} \rightarrow \mathcal{Z}$ be any d_y -private mechanism. Then for all row vectors $\boldsymbol{\mu} \geq \mathbf{0}$ satisfying $\boldsymbol{\mu} \Phi = \mathbf{u}$, and for all priors $\boldsymbol{\pi}$, we have that*

$$\mathcal{L}(H, \boldsymbol{\pi}) \leq \log_2(|\mathcal{Y}| \sum_{y \in \mathcal{Y}} \mu_y)$$

5.3 Quantifying the leakage about the database

In the previous section we considered the information about the query result that is revealed by a mechanism H . This information was measured by the min-mutual information $\mathcal{L}(H, \boldsymbol{\pi})$.

We now turn our attention to the case of standard differential privacy, with the goal of quantifying the information about the *database* that is conveyed by a differentially private mechanism K (not necessarily oblivious). Intuitively, we wish to minimize this information to protect the privacy of the users, contrary to the utility which we aim at maximizing. We can apply the results of the previous section by considering the full mechanism K , mapping databases $\mathcal{V} = V^u$ to outputs (recall that u is the number of individuals in the database and V the universe of values). Differential privacy corresponds to ϵd_h -privacy, where d_h is the *Hamming distance* on the domain \mathcal{V} of databases. Correspondingly ϵd_h -regularity will concern priors $\boldsymbol{\pi}$ on databases \mathcal{V} .

In this case, $\mathcal{L}(K, \boldsymbol{\pi})$ measures the information about the database conveyed by the mechanism, which we refer to as “min-entropy leakage”, and the bounds from the previous section can be directly applied. However, since we now work on a specific metric space $(\mathcal{V}, \epsilon d_h)$, we can obtain a closed expression for the bound of Corollary 2. We start by observing that due to the symmetry of the graph, the uniform prior \mathbf{u} is ϵd_h -regular for all $\epsilon > 0$. More precisely, we can show that the vector $\boldsymbol{\mu}$ of size \mathcal{V} having all elements equal to

$$\left(\frac{e^\epsilon}{|V|(|V| - 1 + e^\epsilon)} \right)^u$$

satisfies $\boldsymbol{\mu} \Phi = \mathbf{u}$ and $\boldsymbol{\mu} \geq \mathbf{0}$. Thus, applying Corollary 2 we get the following result.

Theorem 3 (min-entropy leakage upper bound). *Let $\mathcal{V} = V^u$ be a set of databases, let $\epsilon > 0$, and let K be an ϵ -differentially private mechanism. Then for all priors $\boldsymbol{\pi}$ on \mathcal{V} , we have:*

$$\mathcal{L}(K, \boldsymbol{\pi}) \leq u \log_2 \frac{|V| e^\epsilon}{|V| - 1 + e^\epsilon}$$

This bound determines the maximum amount of information that *any* ϵ -differentially privacy mechanism can leak about the database (independently from the underlying query). The bound was first conjectured in [17] and independently proven in [19]; our technique gives an alternative and arguably more intuitive proof of this result.

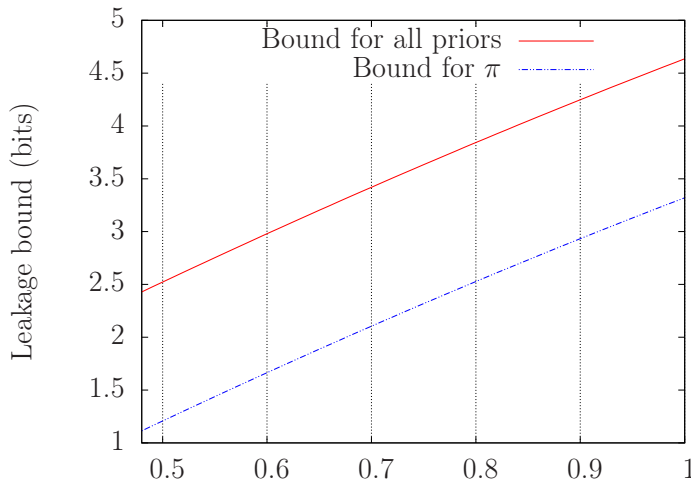


Figure 4: Leakage bounds for various values of ϵ

Note that the above bound holds for *all priors*. If we restrict to a *specific* ϵd_h -regular prior π , then we can get better results by using the bound of Proposition 5 which depends on the actual prior. This is demonstrated in the following example.

Example 3. Consider a database of 5 individuals, each having one of 4 possible values, i.e. $\mathcal{V} = V^u$ with $V = \{1, 2, 3, 4\}$ and $u = 5$. Assume that each individual selects a value independently from the others, but not all values are equally probable; in particular the probabilities of values 1, 2, 3, 4 are 0.3, 0.27, 0.23, 0.2 respectively. Let π be the corresponding prior on \mathcal{V} that models this information. We have numerically verified that for all $0.48 \leq \epsilon \leq 1$ (with step 0.01) π is ϵd_h -regular. Thus we can apply Proposition 5 to get an upper bound of $\mathcal{L}(K, \pi)$ for this prior.

The resulting bound, together with the general bound for all priors from Theorem 3, are shown in Figure 4. We see that restricting to a specific prior provides a significantly better bound for all values of ϵ . For instance, for $\epsilon = 0.5$ we get that $\mathcal{L}(K, \pi) \leq 1.2$ for this π , while $\mathcal{L}(K, \pi) \leq 2.5$ for all priors π .

6 Tight-constraints mechanisms

In general, the bounds for the utility (Theorem 2) and the min-mutual information (Proposition 5) are not tight. That is for a given metric d_y on a domain \mathcal{Y} , there may be no d_y -private mechanism H that meets these bounds. Nevertheless, they provide ultimate limits, induced by the d_y -privacy constraints, for all d_y -private mechanisms and d_y -regular priors. These bounds are simultaneously tight if the condition $\Phi \text{diag}(HR) = \mathbf{1}$ is satisfied (note that this condition is independent of the underlying prior). In this section we exploit this ‘tightness’ condition and investigate the mechanisms that, whenever exist, satisfy this condition and are therefore optimal for the entire region of d_y -regular priors. We call these mechanisms *tight-constraints mechanisms*.

Definition 6 (A tight-constraints mechanism). For a metric d_y , a mechanism $H : \mathcal{Y} \rightarrow \mathcal{Y}$ is called a *tight-constraints* mechanism iff it satisfies the following conditions for all $y, y' \in \mathcal{Y}$.

$$e^{-d_y(y, y')} h_{y' y'} = h_{y y'}. \quad (12)$$

It is important to note that, in general, there may exist zero, one or more tight-constraints mechanisms for a given metric d_y . The above definition enforces $|\mathcal{Y}|(|\mathcal{Y}| - 1)$ linearly independent equations, referred to as the ‘*tight constraints*’. Additionally it must also hold that $\sum_{y' \in \mathcal{Y}} h_{y y'} = 1$ for all $y \in \mathcal{Y}$. Thus we have, in total, $|\mathcal{Y}||\mathcal{Y}|$ equations. If these equations are linearly independent, then they solve to unique values. If these values are non-negative, then they determine a *unique* tight-constraints mechanism. On the other hand, if these equations are not linearly independent, then there may be multiple solutions with non-negative entries, in which case we have multiple tight-constraints mechanisms for d_y .

6.1 Properties

The first feature that follows immediately from the definition of tight-constraints mechanisms, for a metric d_y , is that they satisfy d_y -privacy:

Proposition 6 (d_y -privacy). *For a given metric d_y , every tight-constraints mechanism is d_y -private.*

Proof. For a tight-constraints mechanism \hat{H} , we want to show that for every pair of query results y, y' and every output z , we have

$$\hat{h}_{yz} \leq e^{d_y(y, y')} \cdot \hat{h}_{y'z}. \quad (13)$$

By Definition 6, for every pair of elements y, y' and every output z , we have

$$\hat{h}_{y'z} = e^{-d_y(y', z)} \cdot \hat{h}_{zz} \quad \text{and} \quad \hat{h}_{yz} = e^{-d_y(y, z)} \cdot \hat{h}_{zz}. \quad (14)$$

If $\hat{h}_{zz} = 0$ then $\hat{h}_{y'z} = \hat{h}_{yz} = 0$. In this case, Condition (13) is satisfied. Otherwise (i.e. if $\hat{h}_{zz} \neq 0$), both $\hat{h}_{y'z}$ and \hat{h}_{yz} are non-zero, and it follows from Equations (14) that, for all inputs y and y' , and every output z ,

$$\hat{h}_{y'z} / \hat{h}_{yz} = e^{-(d_y(y', z) - d_y(y, z))}.$$

By the triangle inequality, we have that $d_y(y', z) - d_y(y, z) \leq d_y(y, y')$. Knowing also that $e^{-1} < 1$, it follows from the above inequality that

$$\hat{h}_{y'z} / \hat{h}_{yz} \geq e^{-d_y(y, y')}.$$

The above inequality is equivalent to Condition (13) of d_y -privacy. □ □

Thanks to the above property, we can give a further useful characteristic for the tight-constraints mechanisms distinguishing them from other d_y -private mechanisms. More precisely, the following proposition identifies a linear algebraic condition that is satisfied *only* by the tight-constraints mechanisms for the given metric d_y :

Proposition 7 (diagonal characterization). *For a metric d_y , a d_y -private mechanism $H : \mathcal{Y} \rightarrow \mathcal{Y}$ is a tight-constraints mechanism iff*

$$\Phi \operatorname{diag}(H) = \mathbf{1}. \quad (15)$$

Proof. If H is a tight-constraints mechanism, then by Definition 6 we have that $h_{yy'} = e^{-d_y(y,y')} h_{y'y'}$ for all $y, y' \in \mathcal{Y}$. It also holds that $\sum_{y' \in \mathcal{Y}} h_{yy'} = 1$ for all $y \in \mathcal{Y}$. Combining these equations yields

$$\sum_{y' \in \mathcal{Y}} e^{-d_y(y,y')} h_{y'y'} = 1, \quad \forall y \in \mathcal{Y}. \quad (16)$$

Using the privacy-constraints matrix Φ , the above equations can be written in the matrix form (15). Now we prove the other direction of implication as follows. Suppose that Eq. (16) (which is equivalent to Eq. (15)) is satisfied by a d_y -private mechanism H . Then it holds for all $y, y' \in \mathcal{Y}$ that $h_{yy'} \geq e^{-d_y(y,y')} h_{y'y'}$. Suppose for a contradiction that this inequality is *strict* for some $y, y' \in \mathcal{Y}$, i.e. $h_{yy'} > e^{-d_y(y,y')} h_{y'y'}$. Then $\sum_{y' \in \mathcal{Y}} h_{yy'} > \sum_{y' \in \mathcal{Y}} e^{-d_y(y,y')} h_{y'y'} = 1$, where the last equality holds by Eq. (16). That is, the sum of the entries of a row in H is strictly greater than 1 which violates the validity of H . □ □

The above proposition provides a way to check the existence of, and also compute, the tight-constraints mechanisms for a given metric d_y . Since Condition (15) is satisfied only by these mechanisms, there is at least one tight-constraints mechanism if there is a vector \mathbf{z} , with non-negative entries, that satisfies the equation $\Phi \mathbf{z} = \mathbf{1}$. In this case a tight-constraints mechanism is obtained by setting its diagonal to \mathbf{z} , and evaluating the non-diagonal entries from the diagonal using Eqs. (12).

Now we turn our attention to the region of d_y -regular priors and identify the mechanisms that are optimal with respect to both utility and min-mutual information in this region. Precisely, we show that the set of these optimal mechanism consists exactly of all mechanisms that can be *mapped* to a tight-constraints one using some remap R .

Theorem 4 (Optimality). *Let d_y be a metric for which at least one tight-constraints mechanism exists. Then a d_y -private mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$ is d_y -optimal (wrt both utility and min-mutual information) for every d_y -regular prior π iff there is a remap $R : \mathcal{Z} \rightarrow \mathcal{Y}$ such that HR is a tight-constraints mechanism for d_y .*

Proof. If there exists a tight-constraints mechanism H' for a given metric d_y , then H' must satisfy Eq. (15). This implies that the upper-bound in Theorem 2 is reachable by H' and the identity remap. Thus the upper-bound, in this case, is tight. Now consider a d_y -private mechanism $H : \mathcal{Y} \rightarrow \mathcal{Z}$. By Theorem 2, H meets that upper bound for the utility (and therefore is d_y -optimal) iff it satisfies the condition $\Phi \operatorname{diag}(HR) = \mathbf{1}$, with some remap R . Since H is d_y -private, HR is also d_y -private. Now by Proposition 7, satisfying the condition $\Phi \operatorname{diag}(HR) = \mathbf{1}$ (meaning that H is optimal) is equivalent to that HR is a tight-constraints mechanism (for d_y). Using the relation, given by Proposition 4, between utility and min-mutual information, the same argument holds for the latter. □ □

Observe that tight-constraints mechanisms are optimal because they are mapped to themselves by the identity remap. In the light of Theorem 4, we consider the special case of the uniform prior, denoted by \mathbf{u} , where all results in \mathcal{Y} are equally likely. Note that this prior corresponds to users having unbiased knowledge about the query results, i.e. they assume that all the true results \mathcal{Y} are yielded, by executing the query, with the same probability. Firstly, the following lemma proves an equivalence between the existence of at least one tight-constraints mechanism on one hand and the uniform prior \mathbf{u} being d_y -regular on the other hand.

Proposition 8. *For a given metric d_y , there exists at least one tight-constraints mechanism iff the uniform prior \mathbf{u} is d_y -regular.*

Proof. By Proposition 7, if there is at least a tight-constraints mechanism \hat{H} , then Eq. (15) must hold for this mechanism. Taking the transpose of both sides in this equation, and noting that $\Phi^t = \Phi$ (because Φ is symmetric), then we get that

$$(\text{diag}(\hat{H}))^t \cdot \Phi = \mathbf{1}^t.$$

Scaling the above equation by $1/|\mathcal{Y}|$ yields the row vector \mathbf{u} , the uniform prior, on the right hand side. Thus if a tight-constraints mechanism \hat{H} , exists then

$$(1/|\mathcal{Y}|) (\text{diag}(\hat{H}))^t \cdot \Phi = \mathbf{u}.$$

which means (By Def. 4) that \mathbf{u} is d_y -regular, because the row vector $(\text{diag}(\hat{H}))^t$ has only non-negative entries. For the opposite implication, assume that \mathbf{u} is d_y -regular. Then by the definition there is a row vector $\boldsymbol{\mu}$ with non-negative entries such that $\boldsymbol{\mu} \Phi = \mathbf{u}$. Taking the transpose of both sides, and multiplying by $|\mathcal{Y}|$, yields that Eq. (15) is satisfied for H , whose diagonal is given by $\text{diag}(H) = |\mathcal{Y}| \cdot \boldsymbol{\mu}^t$ (non-negative). Thus there exists a tight-constraints mechanism for d_y . \square \square

It is worth noticing that in general the region of d_y -regular priors may or may not include the uniform prior. However, as shown earlier in Section 4, this region is enlarged and converges to the entire prior space as the distances $d_y(y, y') \rightarrow \infty$ for all $y \neq y'$. In particular the d_y -regular priors accommodate the uniform prior \mathbf{u} if d_y is scaled up by an appropriate factor.

In the case of ϵ -differential privacy it holds that $d_y = \epsilon d_h$ where d_h is the Hamming distance on databases. Thus there is always a threshold ϵ^* , above which the uniform prior \mathbf{u} is ϵd_h -regular. This can provide a design criteria to *select* a setting for ϵ such that, according to Proposition 8, there is a tight-constraints mechanism that is optimal for all ϵd_h -regular priors.

Using Proposition 8, we can describe the optimal mechanisms for the uniform prior as a corollary of Theorem 4.

Corollary 3. *Let d_y be a metric for which there exists at least one tight-constraints mechanism. Then a mechanism H is d_y -optimal for the uniform prior on \mathcal{Y} iff HR is a tight-constraints mechanism for some remap $R : \mathcal{Z} \rightarrow \mathcal{Y}$.*

In summary, the existence of tight-constraints mechanisms and their structures depend on the given metric. The choice of such metric corresponds to the required privacy guarantee. Consider

in particular the conventional ϵ -differential privacy, where any two *adjacent* elements in a domain \mathcal{Y} are required to be indistinguishable relative to ϵ . In this case, the domain \mathcal{Y} and its adjacency relation \sim_f are modeled by the graph $G = (\mathcal{Y}, \sim_f)$; and the requirement of satisfying ϵ -differential privacy for \mathcal{Y} translates in our general model to the metric $d_y(y, y') = \epsilon d_{\sim_f}(y, y')$, where $d_{\sim_f}(y, y')$ is the graph distance between y, y' . With this metric, we find that tight-constraints mechanisms capture other known differentially-private mechanisms. For example, if we set \mathcal{Y} to be the output domain of a counting query executed on a database, we find that the tight-constraints mechanism for \mathcal{Y} is exactly the *truncated-geometric mechanism*, which was shown by [13] to be optimal for every prior. Also, we instantiate, in the following, the tight-constraints mechanism when the metric space (\mathcal{Y}, d_y) satisfies a certain symmetry. This symmetry captures, in particular, the graphs for which an optimal mechanism is constructed in [19] for the uniform prior \mathbf{u} . Once again this mechanism is precisely a tight-constraints one. Note that an additional conclusion which we add here is that this mechanism is optimal not only for \mathbf{u} but also for all d_y -regular priors.

6.2 Tight-constraints mechanism for symmetric metric spaces

We consider the mechanisms that satisfy d_y -privacy for a given domain \mathcal{Y} . We focus here on the metric spaces (\mathcal{Y}, d_y) that satisfy a certain symmetry which we call *ball-size symmetry*. To describe this property, we recall the standard notion of balls in metric spaces: a *ball* of radius r around a point $y \in \mathcal{Y}$ is the set $B_r^{d_y}(y) = \{y' \in \mathcal{Y} : d_y(y, y') \leq r\}$. Now we define the ball-size symmetry as follows.

Definition 7 (ball-size symmetry). A metric space (\mathcal{Y}, d_y) is said to be ball-size symmetric if for all $y, y' \in \mathcal{Y}$, and all radii r , we have $|B_r^{d_y}(y)| = |B_r^{d_y}(y')|$.

Note that the above condition is equivalent to saying that for any $y \in \mathcal{Y}$, the number of elements that are at distance r from y depends only on r , allowing us to write this number as n_r . Inspecting the privacy-constraints matrix Φ in this case, we observe that the row sum $\sum_{y'} \phi_{yy'}$ for every $y \in \mathcal{Y}$ is the same and equal to $\sum_r n_r e^{-r}$. This means that the column vector \mathbf{z} , of which every element is equal to $1/\sum_r n_r e^{-r}$, satisfies $\Phi \mathbf{z} = \mathbf{1}$ and therefore yields (by Proposition 7) the diagonal of a tight-constraints mechanism H . The other (non-diagonal) entries of H follow from the diagonal as in Definition 6. Thus we conclude the following result.

Proposition 9 (tight-constraints mechanism for symmetric metric spaces). *For any metric space (\mathcal{Y}, d_y) satisfying ball-size symmetry there is a tight-constraints mechanism $H : \mathcal{Y} \rightarrow \mathcal{Y}$ which is given as $h_{yy'} = e^{d_y(y, y')} / \sum_r n_r e^{-r}$.*

The main consequence of the above proposition is that the mechanism H is optimal for every d_y -regular prior including the uniform prior \mathbf{u} .

The above result generalizes and extends a result by [19] in the context of differential privacy. The authors of [19] considered two types of graphs: distance-regular and vertex-transitive graphs. They constructed for these graphs an ϵ -differentially private mechanism optimal for the uniform prior. As shown earlier ϵ -differential privacy for a graph (\mathcal{Y}, \sim_f) translates in our setting to the




Figure 5: Adjacency graphs

metric space $(\mathcal{Y}, \epsilon d_{\sim_f})$. It can be easily seen that if (\mathcal{Y}, \sim_f) is either distance-regular or vertex-transitive, the corresponding metric space $(\mathcal{Y}, \epsilon d_{\sim_f})$ is ball-size symmetric. Therefore, we can instantiate the tight-constraints mechanism of Proposition 9 to ϵd_{\sim_f} , which gives exactly the optimal mechanism constructed in [19]. Hence, we directly obtain the same optimality results, and moreover our analysis shows that this mechanism is optimal on the entire region of ϵd_{\sim_f} -regular priors, instead of only the uniform one.

7 Case-studies

In this section we show the usefulness of the tight-constraints mechanism by applying it to two contexts: standard differential privacy and geo-indistinguishability.


7.1 Differential privacy: sum and 2-count queries

We evaluate the tight constraints mechanism for two families of queries, namely sum and 2-count queries. For each family, we apply the mechanism on databases consisting of u individuals each having an integer value between 0 and v , and we compare its utility to the geometric mechanism.

It is well-known that no universally optimal mechanism exists for these families; in particular, the geometric mechanism, known to be optimal for a single counting query, is not guaranteed to be optimal for sum queries or multiple counting queries. On the other hand, as discussed in the previous section, tight-constraints mechanisms, whenever they exist, are guaranteed to be optimal within the region of regular priors.

The comparison is made as follows: for each query, we numerically compute the smallest ϵ (using a step of 0.01) for which a tight-constraints mechanism exists (i.e. for which the uniform prior \mathbf{u} is ϵd_{\sim_f} -regular, see Proposition 8). Then we compute the utility (using an optimal remap) of both the tight constraints and the geometric mechanisms, for a range of ϵ starting from the minimum one. Note that the tight constraint mechanism exists for any ϵ greater than the minimum one.

Sum query Let f be the query returning the sum of the values for all individuals, thus it has range $\mathcal{Y} = \{0, \dots, vu\}$. By modifying the value of a single individual, the outcome of the query

Figure 6: Utility for various values of ϵ

can be altered by at most v (when changing the value from 0 to v), thus two elements $i, j \in \mathcal{Y}$ are adjacent iff $|i - j| \leq v$. The induced graph structure on \mathcal{Y} is shown in Figure 5(a) (for the case $v = 3$).

For our case-study we numerically evaluate this query for $u = 150, v = 5$ and for the uniform prior. We found that the minimum ϵ for which a tight-constraints mechanism exists (and is in fact unique since Φ is invertible) is 0.8. Figure 6(a) shows the utility of the tight-constraint mechanism, as well as that of the geometric mechanism, for values of ϵ between 0.8 and 1.3, the uniform prior and using and optimal remap. We see that the tight-constraints mechanism provides significantly higher utility than the geometric mechanism in this case.

2-count query Consider now the query f consisting of 2 counting queries (i.e. reporting the number of users satisfying properties p_1 and p_2), thus it has range $\mathcal{Y} = \{0, \dots, u\} \times \{0, \dots, u\}$. By modifying the value of a single individual, the outcome of each counting query can be altered by at most 1, thus two answers $(i_1, i_2), (j_1, j_2) \in \mathcal{Y}$ are adjacent iff $|i_1 - j_1| \leq 1$ and $|i_2 - j_2| \leq 1$. The induced graph structure on \mathcal{Y} is shown in Figure 5(b).

We evaluate this query for $u = 30$ and for the uniform prior. We found that the minimum ϵ for which a tight-constraints mechanism exists is 0.9. Figure 6(b) shows the utility of the two mechanisms (with the geometric being applied independently to each counting query) for values of ϵ between 0.9 and 1.3 and the uniform prior. Similarly to the sum query, we see that the tight-constraints mechanism provides significantly higher utility than the geometric mechanism in this case.

7.2 Geo-indistinguishability

As discussed in Section 3.3, geo-indistinguishability is a notion of location privacy obtained by taking $d_x = \epsilon d_2$, where d_2 is the Euclidean distance between locations. In [9] it is shown that a planar version of the Laplace mechanism satisfies ϵ -geo-indistinguishability. The Planar Laplace

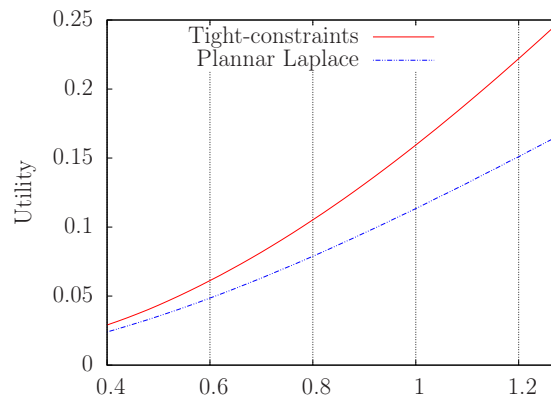


Figure 7: Utility of location privacy mechanisms for various values of ϵ

mechanism is continuous, having as input and output the full \mathbb{R}^2 , but in the case of a finite number of locations it can be discretized and truncated while still satisfying geo-indistinguishability (for a slightly adjusted ϵ).

Although the Planar Laplace mechanism is simple, efficient and easy to implement, it provides no optimality guarantees. On the other hand, for any finite number of locations, the tight-constraints mechanism, if it exists, is guaranteed to be optimal for ϵd_2 -regular priors. In this section we compare the two mechanisms on a grid of 100×100 locations, with step size 1 km.

Note that constructing the tight-constraints mechanism involves inverting the matrix Φ , which can be done in time $O(|\mathcal{X}|^{2.376})$ using the Coppersmith-Winograd algorithm. This complexity is much lower than that of recent methods for computing optimal location obfuscation mechanisms. For instance, the well-known method of Shokri et al. [23] – which uses the adversary’s expected error as the metric of privacy – involves solving large linear optimization problems and was evaluated to a grid of only 30 locations (compared to the 10,000 locations in our grid).

Figure 7 shows the utility of the two mechanisms for ϵ ranging from 0.4 to 1.3 and for a uniform prior. As expected, the tight-constraints mechanism offers significantly higher utility than the Planar Laplace mechanism for the same ϵ .

It should be emphasized, however, that our optimality results hold for the binary gain function, which corresponds to an attacker trying to guess the true location of the user (the utility being the probability of a correct guess). This might often be meaningful, especially when the grid size is big: guessing any incorrect cell could be considered equally bad. But it is also common to consider gain functions taking the distance between locations into account, with respect to which the tight-constraints mechanism is not guaranteed to be optimal.

8 Conclusion and future work

In this paper we have continued the line of research initiated by [13, 14] about the existence of differentially-private mechanisms that are universally optimal, i.e., optimal for all priors. While the positive result of [13] (for counting queries) and the negative one of [14] (for essentially

all other queries) answer the question completely, the latter sets a rather dissatisfactory scenario, since counting queries are a very specific kind of queries, and in general users can be interested in very different queries. We have then considered the question whether we can achieve optimality with the same mechanism for a restricted class of priors. Fortunately the answer is positive: we have identified a region of priors, called d_y -regular, and a mechanism, called tight-constraints, which is optimal for all the priors in this region. We have also provided a complete and effectively checkable characterization of the conditions under which such mechanism exists, and an effective method to construct it. As a side result, we have improved on the existing bounds for the min-entropy leakage induced by differential privacy. More precisely, we have been able to give specific and tight bounds for each d_y -regular prior, in general smaller than the bound existing in the literature for the worst-case leakage (achieved by the uniform prior [18]).

So far we have been studying only the case of utility for binary gain functions. In the future we aim at lifting this limitation, i.e. we would like to consider also other kinds of gain. Furthermore, we intend to study how the utility decreases when we use a tight-constraints mechanism outside the class of d_y -regular priors. In particular, we aim at identifying a class of priors, larger than the d_y -regular ones, for which the tight-constraints mechanism is close to be optimal.

Bibliography

- [1] Sabelfeld, A., Myers, A.C.: Language-based information-flow security. *IEEE Journal on Selected Areas in Communications* **21**(1) (2003) 5–19
- [2] Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: Anonymity protocols as noisy channels. *Inf. and Comp.* **206**(2–4) (2008) 378–401
- [3] Dwork, C.: A firm foundation for private data analysis. *Communications of the ACM* **54**(1) (2011) 86–96
- [4] Dwork, C.: Differential privacy. In: *Proc. of ICALP*. Volume 4052 of LNCS., Springer (2006) 1–12
- [5] Dwork, C., Mcsherry, F., Nissim, K., Smith, A.: Calibrating noise to sensitivity in private data analysis. In: *Proc. of TCC*. Volume 3876 of LNCS., Springer (2006) 265–284
- [6] Narayanan, A., Shmatikov, V.: Robust de-anonymization of large sparse datasets. In: *Proc. of S&P*. (2008) 111–125
- [7] Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: *Proc. of S&P*, IEEE (2009) 173–187
- [8] Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the scope of Differential Privacy using metrics. In: *Proc. of PETS*. Volume 7981 of LNCS., Springer (2013) 82–102

- [9] Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K., Palamidessi, C.: Geo-indistinguishability: differential privacy for location-based systems. In: Proc. of CCS, ACM (2013) 901–914
- [10] Gaboardi, M., Haeberlen, A., Hsu, J., Narayan, A., Pierce, B.C.: Linear dependent types for differential privacy. In: Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. POPL '13, New York, NY, USA, ACM (2013) 357–370
- [11] Barthe, G., Olmedo, F.: Beyond differential privacy: Composition theorems and relational logic for f-divergences between probabilistic programs. In: Automata, Languages, and Programming - 40th Int. Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part II. Volume 7966 of LNCS., Springer (2013) 49–60
- [12] Barthe, G., Köpf, B., Olmedo, F., Béguelin, S.Z.: Probabilistic relational reasoning for differential privacy. *ACM Trans. Program. Lang. Syst* **35**(3) (2013) 9
- [13] Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. In: Proc. of STOC, ACM (2009) 351–360
- [14] Brenner, H., Nissim, K.: Impossibility of differentially private universally optimal mechanisms. In: Proc. of FOCS, IEEE (2010) 71–80
- [15] Smith, G.: On the foundations of quantitative information flow. In: Proc. of FOSSACS. Volume 5504 of LNCS., Springer (2009) 288–302
- [16] Braun, C., Chatzikokolakis, K., Palamidessi, C.: Quantitative notions of leakage for one-try attacks. In: Proc. of MFPS. Volume 249 of ENTCS., Elsevier (2009) 75–91
- [17] Barthe, G., Köpf, B.: Information-theoretic bounds for differentially private mechanisms. In: Proc. of CSF, IEEE (2011) 191–204
- [18] Alvim, M.S., Andrés, M.E., Chatzikokolakis, K., Degano, P., Palamidessi, C.: Differential Privacy: on the trade-off between Utility and Information Leakage. In: Postproceedings of the 8th Int. Workshop on Formal Aspects in Security and Trust (FAST). Volume 7140 of LNCS., Springer (2011) 39–54
- [19] Alvim, M.S., Andrés, M.E., Chatzikokolakis, K., Palamidessi, C.: On the relation between Differential Privacy and Quantitative Information Flow. In: Proc. of ICALP. Volume 6756 of LNCS., Springer (2011) 60–76
- [20] Ardagna, C.A., Cremonini, M., Damiani, E., di Vimercati, S.D.C., Samarati, P.: Location privacy protection through obfuscation-based techniques. In: Proc. of DAS. Volume 4602 of LNCS., Springer (2007) 47–60
- [21] Shokri, R., Theodorakopoulos, G., Boudec, J.Y.L., Hubaux, J.P.: Quantifying location privacy. In: Proc. of S&P, IEEE (2011) 247–262

- [22] Rényi, A.: On Measures of Entropy and Information. In: Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability. (1961) 547–561
- [23] Shokri, R., Theodorakopoulos, G., Troncoso, C., Hubaux, J.P., Boudec, J.Y.L.: Protecting location privacy: optimal strategy against localization attacks. In: Proc. of CCS, ACM (2012) 617–627

MEALS Partner Abbreviations

SAU: Saarland University, D

RWT: RWTH Aachen University, D

TUD: Technische Universität Dresden, D

INR: Institut National de Recherche en Informatique et en Automatique, FR

IMP: Imperial College of Science, Technology and Medicine, UK

ULEIC: University of Leicester, UK

TUE: Technische Universiteit Eindhoven, NL

UNC: Universidad Nacional de Córdoba, AR

UBA: Universidad de Buenos Aires, AR

UNR: Universidad Nacional de Río Cuarto, AR

ITBA: Instituto Tecnológico Buenos Aires, AR